CDMA/HDR: A Bandwidth-Efficient High-Speed Wireless Data Service for Nomadic Users

Paul Bender, Peter Black, Matthew Grob, Roberto Padovani, Nagabhushana Sindhushayana, and Andrew Viterbi, QUALCOMM, Incorporated

ABSTRACT

This article presents an approach to providing very high-data-rate downstream Internet access by nomadic users within the current CDMA physical layer architecture. Means for considerably increasing throughput by optimizing packet data protocols and by other network and coding techniques are presented and supported by simulations and laboratory measurements. The network architecture, based on Internet protocols adapted to the mobile environment, is described, followed by a brief discussion of economic considerations in comparison to cable and DSL services.

INTRODUCTION

The rapid growth and nearly universal coverage of industrialized nations and regions by digital wireless telephony gives rise to an increasing demand for data services as well. While current offerings are for data rates equivalent to those provided by wireline modems a decade or more ago, the gap is closing. Standards are already approved and chip sets are available for providing data rates above 64 kb/s within this calendar year (and century). Just beyond this horizon, however, service providers are already planning for wireless data rates above 2 Mb/s, approaching those of wireline, digital subscriber line (DSL), and cable. Whether such a wireless service can be made technically and economically competitive with wireline and cable is not the main issue, although we shall address this briefly in the last section. What will drive such a service is the demand for rapid low-latency availability of Internet access to nomadic users. In the next section we describe the characteristics and perceived needs of this user class. We then proceed to explore the characteristics of data requirements for speed and latency, after which we present a technical system solution tailored to these requirements and the characteristics of a specific implementation as an evolution of existing CDMA base station and subscriber terminal architectures. In the final section we briefly discuss the economics of such a system deployment.

CHARACTERISTICS OF NOMADIC USER DATA DEMAND

In business and the professions, the individual is often absent from her or his normal workplace. To continue to be productive on the road, both in transit and at business or professional meetings, connectivity to data at one's principal workplace and more broadly to other databases accessible through the Web is essential. Generally, members of this nomadic user class demand the same data service normally available in their home base. Often cited examples of the nature of such services are e-mail retrieval, Web browsing, ordering airline tickets, hotel reservations, obtaining stock quotes, and report retrieval, in such locations as airport lounges, hotel rooms, and meeting places, in each case without recourse to the limited or interface unfriendly facilities available in such places. In fact, one need not necessarily look beyond corporate boundaries; professional employees often spend nearly as much time in company conference rooms as in their own offices, and rarely are such rooms equipped with the number of ports needed to connect the majority of participants' laptops.[1]

The nature of such data traffic is decidedly asymmetric. A much higher forward (or downlink) rate is required from the access point (base station) than that generated by the access terminal (user terminal) in the reverse (uplink) direction. Furthermore, just as does the fixed user, the nomadic user expects a response to her or his request which does not suffer from excessive latency. Our goal is to satisfy these needs with an evolutional approach which minimizes the time and cost for providing such capabilities in existing cellular infrastructure and with terminals that differ only at digital baseband from existing cellular and personal communications systems (PCS) handsets.

[1] *The obvious alternative of private campus wireless systems will be discussed in the last section.*

0163-6804/00/$10.00 © 2000 IEEE

THE TECHNICAL CHALLENGE

Little information has been gathered on the exigencies of nomadic users and the networks to serve them, simply because, with the exception of a very small percentage of low-speed data service, such networks have not existed. On the other hand, with digital cellular networks in place for nearly a decade and with large numbers of mobile users served for several years, a great deal is known about the characteristics of digital wireless networks and their mobile users. A major step in the perfection of digital cellular technology was the development and standardization of code-division multiple access (CDMA) wireless systems and their adoption by the majority of North American, Korean, and Japanese carriers and manufacturers. Having proven its superiority to other access techniques, it is now being imitated (sometimes in a modified form) by most of the carriers and manufacturers who were the initial holdouts and skeptics of its viability.

CDMA was designed for efficient reverse (uplink) and forward (downlink) operations. It was initially widely believed that the reverse direction in which multiple users access each base station, hence representing multiple sources of interference with one another, would be the capacity limiting (or bottleneck) direction. This assumption turned out to be incorrect; the forward (downlink) was the initial bottleneck for three principal reasons:

- Interference on the reverse link enjoys the advantage of the law of large numbers, whereby the cumulative interference from multiple low-power transmitters tends to be statistically stable. The forward link, on the other hand, suffers interference from a small number of other high-power base stations. This becomes particularly serious at the vertices of the (imaginary) cellular hexagon where the transmitting base station and two other interfering base stations are equidistant from the intended user. This situation is relieved by soft handoff, where two or more base stations transmit to the user simultaneously.
- But soft handoff, while greatly diminishing interference, which itself increases capacity, still overall diminishes the forward link[2] capacity because an additional CDMA carrier must be assigned in the newly added base station. Depending on the region of (or criterion for) soft handoff, this can cause greater or lesser reduction.
- While on the reverse link, fast and accurate power control of multiple users is evidently critical to operation and capacity realization, it was initially felt in producing the first CDMA standard, cdmaOne (IS-95-A [1]), that forward link power control could be much slower. This turned out to reduce forward link capacity.

The second and third limiting causes have been eliminated or considerably diminished in the evolutionary revisions of cdmaOne (IS-95-B and CDMA2000). Fast power control is now implemented in the forward link, and the region of (criterion for) soft handoff has been diminished. The first cause (sometimes called the "law of small numbers"), remains, however. These improvements have brought the forward (downlink) capacity to parity with the reverse (uplink) capacity. But for high-speed data, such as downloading from the Internet, this is *not enough*. The downlink demand is likely to be several times greater than the uplink. The rest of this article deals with new approaches which will further increase the downlink capacity by a factor of three to four for data applications only.[3]

THE TECHNICAL APPROACH TO HIGH-SPEED DATA

Most data applications differ fundamentally from speech requirements in two respects already noted, traffic asymmetry and tolerance to latency. Two-way conversational speech requires strict adherence to symmetry; also, latencies above 100 ms (which corresponds to about 1 kb of data for most speech vocoders) are intolerable. For high-speed data downlinked at 1 Mb/s, for example, 100 ms represents 100 kb or 12.5 kbytes; furthermore, latencies of 10 s are hardly noticeable, and this corresponds to a record of 1.25 Mbytes. Thus, smoothing over a variety of conditions, which is always advantageous for capacity, is easily accomplished.

All communication systems, wired as well as wireless, are greatly improved by a combination of techniques based on three principles:

- Channel measurement
- Channel control
- Interference suppression and mitigation

Our approach employs all three. First, on the basis of the received common pilot from each access point (or base station), each access terminal (subscriber terminal) can measure the received signal-to-noise-plus-interference ratio (SNR). The data rate which can be supported to each user is proportional to its received SNR. This may change continuously, especially for mobile users. Thus, over each user's reverse (uplink) channel, the SNR or equivalently the supportable data rate value is transmitted to the base station. In fact, since typically two or more base stations may be simultaneously tracked, the user indicates the highest among its received SNRs and the identity of the base station from which it is receiving it, and this may need to be repeated frequently (possibly every slot[4]). In this way the downlink channel is controlled as well as measured. Furthermore, by selecting only the best base station, in terms of SNR, to transmit to the user, interference to users of other base stations is reduced. Additionally, since data can tolerate considerably more delay than voice, error-correcting coding techniques which involve greater delay, specifically turbo codes, can be employed which will operate well at lower E_b/N_0, and hence lower SNR and higher interference levels.

Next, we show how unequal latency, for users of disparate SNR levels, can be used to increase throughput. Suppose we can separate users into N classes according to their SNR levels, and corresponding instantaneous rate levels supportable. Thus, user class n can receive slots at rate R_n b/s, where $n = 1,2..N$, and suppose the relative frequency of user packets of class n is P_n.

All communication systems, wired as well as wireless, are greatly improved by a combination of techniques based on three principles: channel measurement, channel control, and interference suppression and mitigation Our approach employs all three.

[2] *Although not the capacity of the reverse link, which soft handoff actually increases.*

[3] *Clearly voice is fundamentally a symmetric service, with stricter latency requirements, as we shall note below.*

[4] *In speech-oriented CDMA, voice frames are 20 ms long. In the next section, we shall establish corresponding lengths for data, which will be called slots. Multiplying R_{av} of (1) by slots per second yields throughput in bits per second.*

Data rate (kb/s)	Packet length (bytes)	FEC rate (b/sym)	Modulation
38.4	128	1/4	QPSK
76.8	128	1/4	QPSK
102.6	128	1/4	QPSK
153.6	128	1/4	QPSK
204.8	128	1/4	QPSK
307.2	128	1/4	QPSK
614.4	128	1/4	QPSK
921.6	192	3/8	QPSK
1228.8	256	1/2	QPSK
1843.2	384	1/2	8PSK
2457.6	512	1/2	16QAM

Table 1. *Various data rates.*

Suppose slots are assigned one at a time successively to each user class. Then the average rate, which we define as throughput, is

$$R_{av} = \sum_{n=1}^{N} P_n R_n \ \text{b/s}. \tag{1}$$

This, of course, means that lower-data-rate (and SNR) users will have proportionately higher latency. For if B bits are to be transmitted altogether for each class, the number of slots (and hence time) required for user class n will be B/R_n, and hence the latency L_n is inversely proportional to R_n.

Suppose, on the other hand, that we require all users to have essentially the same latency[5] irrespective of the R_n they can support. Then as each user class is served, it will be allocated a number of slots inversely proportional to its rate. Let F_n be the number of slots allocated to class n, where $F_n = k/R_n$, k being a constant. In this case, the average rate or throughput is

$$R'_{av} = \frac{\sum_{n=1}^{N} P_n R_n F_n}{\sum_{n=1}^{N} P_n F_n} = \frac{1}{\sum_{n=1}^{N} P_n / R_n} \ \text{b/s}. \tag{2}$$

In this case, however, the latency of all user classes will be the same (assuming the total number of bits K to be large and thus ignoring edge effects).

To assess the cost in throughput for equalizing latency, consider the extreme case of only two user classes, each equally probable ($P_1 = P_2 = 1/2$) but capable of supporting very disparate rates $R_1 = 16$ kb/s, $R_2 = 64R_1 = 1{,}024$ kb/s. Then in the first case, $R_{av} = 520$ kb/s, but $L_1/L_2 = 64$. In the second case, $L_1 = L_2$, but $R'_{av} = 31.51$ kb/s.

To see that there is a more rational allocation which is less "unfair" than a latency ratio of 64, and still achieves a better throughput than R'_{av}, consider a compromise which guarantees that the highest latency is no more than, for example, 8 times the lowest latency. Then in the second case, we would assign 8 slots to class 1 for every slot assigned to class 2. The result would be $L_1/L_2 = 8$ as required and

$$R_{av}'' = (8FP_1R_1 + FP_2R_2)/(FP_1 + FP_2) = 128 \text{ kb/s}.$$

For the general case of N classes and latency ratio L_{max}/L_{min}, it can be shown that the maximum achievable throughput, denoted by C, is

$$C = \frac{\sum_{n=1}^{n_o} P_n + \sum_{n=n_o+1}^{N} P_n (L_{min}/L_{max})}{\sum_{n=1}^{n_o} P_n / R_n + \sum_{n=n_o+1}^{N} (P_n/R_n)(L_{min}/L_{max})} \ \text{b/s}, \tag{3}$$

where $R_1 < R_2 \ldots R_N$ and n_o is such that $R_n \le C$ for all $n \le n_o$, while $R_n > C$ for all $n > n_o$.

Surprisingly, with this maximizing strategy, each user's latency is either L_{max} (for those for which $R_n < C$) or L_{min} (for those for which $R_n > C$). To determine the maximum throughput it is necessary to have a histogram of the achievable rates for users of the wireless network in question. This will be discussed in the next section. Also, as we shall find there, practical numerology considerations may require us to deviate from this strict bimodal latency allocation, although the ratio L_{max}/L_{min} will remain as the principal constraint.

Implementation of High-Data-Rate Code-Division Multiple Access

In the last section we discussed the key factors and parameters of a wireless system designed to optimize the transport of packet data. In the following we will describe such a system design, beginning first with a description of the air interface, to continue in the next section with a description of the network architecture. The design leverages in many ways the lessons learned from the development and operation of CDMA IS-95 networks, but makes no compromises in optimizing the air interface for data services. Furthermore, a compelling economic argument can be made for a design that can reuse large portions (to be exact, all but the baseband signal processing elements) of components and designs already implemented in IS-95 products, both in the access terminals and access points (APs).

Due to the highly asymmetric nature of the service offered, we will focus most of our attention on the downlink. In the IS-95 downlink, a multitude of low-data-rate channels are multiplexed together (with transmissions made orthogonal in the code domain) and share the available base station transmitted power with some form of power control. This is an optimal choice for many low-rate channels sharing a common bandwidth. The situation becomes less optimal when a low number of high-rate users share the channel. The inefficiencies increase further when the same bandwidth is shared between low-rate voice and high-rate data users, since their requirements are vastly different, as discussed previously. It should be noted that

[5] *This is the case for voice. The only difference is that in speech, transmitter power levels are controlled to equalize received power, while here time, in terms of frames, is controlled to equalize energies.*

increasing the bandwidth available for transmission cannot help in this regard if the data rate of the users is increased proportionally as well.

Therefore, a first fundamental design choice is to separate the services, that is, low-rate data (voice being the primary service in this category) from high-rate data services, by using possibly adjacent but nonoverlapping spectrum allocations. To summarize, a better system is one that uses an IS-95 or cdma2000-1X RF carrier to carry voice and a separate high-data-rate (HDR) RF carrier to deliver high-rate packet bursts.

With a dedicated RF carrier, the HDR downlink takes on a different form than that of the IS-95 designs. As shown in Fig. 1, the downlink packet transmissions are time multiplexed and transmitted at the full power available to the AP, but with data rates and slot lengths that vary according to the user channel conditions. Furthermore, when users' queues are empty, the only transmissions from the AP are those of short pilot bursts and periodic transmissions of control information, effectively eliminating interference from idling sectors.

The pilot bursts provide the access terminals with means to accurately and rapidly estimate the channel conditions. Among other parameters, the access terminal estimates the received E_c/N_t of all resolvable multipath components and forms a prediction of the effective received[6] SNR. The value of the SNR is then mapped to a value representing the maximum data rate such a SNR can support for a given level of error performance. This channel state information, in the form of a data rate request, is then fed back to the AP via the reverse link data rate request channel (DRC) and updated as fast as every 1.67 ms, as shown in Fig. 2. The reverse link data request is a 4-bit value that maps the predicted SNR into one of the data rate modes of Table 1. In addition, the access terminal requests transmission from only one sector (that with the highest received SNR) among those comprising the *active set*. Here the definition of active set is identical to that for IS-95 systems, but unlike IS-95, only one sector transmits to any specific access terminal at any given time.

The main coding and modulation parameters are summarized in Table 1.

The forward error correcting (FEC) scheme employs serial concatenated coding and iterative decoding, with puncturing for some of the higher code rates [2].

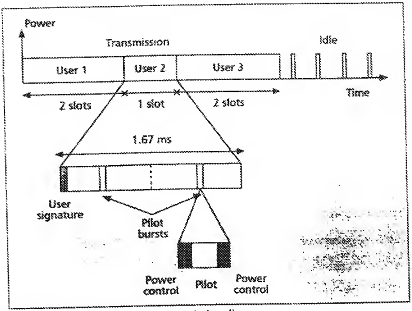

■ Figure 1. *An access point transmission diagram.*

Following the encoder, these traditional signal processing steps are applied: symbol repetition is performed on the lower-data-rate modes: scrambling, channel interleaving, and the appropriate modulation is applied to obtain a constant modulation rate of 1.2288 MHz for all modes. The in-phase and quadrature channels are then each demultiplexed into 16 streams, each at 76.8 kHz, and 16-ary orthogonal covers are applied to each stream. The resulting signal, obtained by adding the 16 data streams, is then spread by quadrature pseudonoise (PN) sequences, bandlimited and upconverted. The resulting RF signal has the same characteristics as an IS-95 signal, thus allowing the reuse of all analog and RF designs developed for IS-95 base stations, including the power amplifiers, and the receiver designs for subscriber terminals.

Table 2 summarizes the SNR required to achieve a 1 percent packet error rate (PER).

Note that at the lower rates this corresponds to E_b/N_0 = 2.5 dB, a result of using iterative decoding techniques on serial concatenated codes, while for the two highest rates, E_b/N_0 increases considerably because 8-phase shift keying (PSK) modulation and 16-quadrature amplitude modulation (QAM) are employed. These were obtained both by bit-exact simulation and

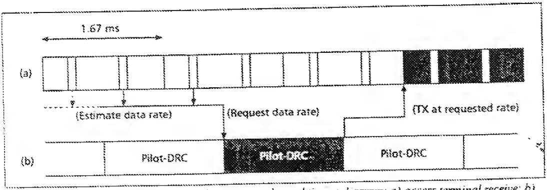

■ Figure 2. *A channel estimation and data request channel timing diagram: a) access terminal receive; b) access terminal transmit.*

[6] *E_c represents the received signal energy density and N_t represents the total nonorthogonal single sided noise density. N_t comprises intercell interference, thermal noise, and possibly nonorthogonal intracell interference.*

corroborated by laboratory measurements with a complete RF link.

At this point we are able to estimate the maximum achievable throughput per sector as discussed in the previous section. Figure 3a shows a graph of the cumulative distribution function of the SNR for a typical embedded sector of a large three-sector network deployed with a frequency reuse of one. In particular, the SNR values are those of the best serving sector and representative of a uniform distribution of users across the coverage area. From the results of Fig. 3a and the knowledge of the SNR required to support a given data rate (Table 1), it is straightforward to derive the histogram of data rates achievable in such an embedded sector. The result is shown in Fig. 3b where the SNRs used in the calculation are those of Table 1 with an additional 2 dB of margin to account for various losses. Finally, Fig. 3c shows the realized throughput per sector per 1.25 MHz of bandwidth versus the parameter L_{max}/L_{min}. Note that the throughput is doubled for a latency ratio $L_{max}/L_{min} = 8$.

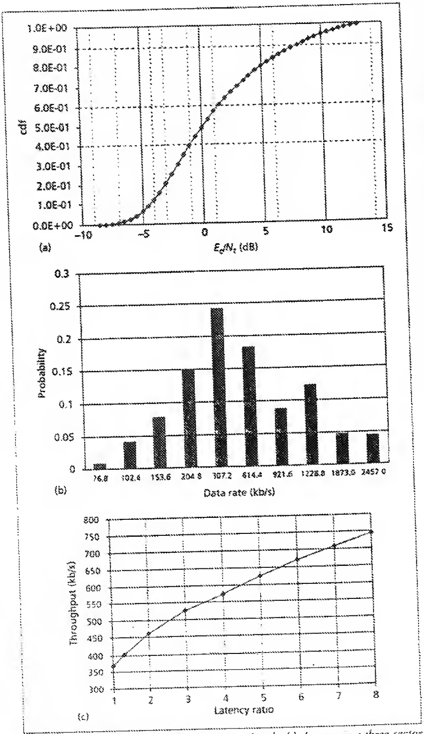

■ Figure 3. *a) E_c/N_t distribution for a typical embedded sector in a three-sector network with universal frequency reuse in each cell; b) data rate histogram; c) sector throughput vs. latency ratio L_{max}/L_{min}.*

Network Architecture

Since the radio link has been designed to provide efficient access to packet data networks, it is natural to turn to the most ubiquitous packet data network — the Internet — when selecting the network architecture. Adopting Internet protocols in the communication between the access terminal and the access network allows users to access the widest variety of information and services, including e-mail, private intranets, and the World Wide Web. Furthermore, the selection of Internet protocols in the design of the access network allows the access network equipment to take advantage of the ever decreasing costs and increasing performance of Internet equipment.

First, we examine the communication link between the access terminal and the access network. Figure 4a shows the protocol stack used in such a link.

In order to carry traffic between the user and the network, we need to select a network-layer protocol. We chose the Internet Protocol (IP) [3] because it is the network-layer protocol of the Internet. The Internet carries its network-layer protocol over a variety of transports. For example, asynchronous transfer mode (ATM) often carries Internet traffic on the Internet backbone, Ethernet often carries Internet traffic on local area networks (LANs), and the Point-to-Point Protocol (PPP) [4] often carries Internet traffic over dialup connections. We chose PPP for the following reasons. First, PPP is widely supported. Moreover, PPP allows the transport of a variety of network-layer protocols, supports methods for

Data rate (kb/s)	E_c/N_t (dB)
38.4	-12.5
76.8	-9.5
102.6	-8.5
153.6	-6.5
204.8	-5.7
307.2	-4.0
614.4	-1.0
921.6	1.3
1228.8	3.0
1843.2	7.2
2457.6	9.5

■ Table 2. *SNR for a 1 percent packet error rate.*

differing quality of service (QoS) requirements, and also supports methods for authentication. Lastly, PPP has low overhead, an important feature for a wireless transport.

It is well known that the Internet carries different types of traffic with different QoS requirements. Some traffic, such as Transmission Control Protocol (TCP) [5] traffic, tends to be more sensitive to errors and less sensitive to delay. Other traffic, such as Real-Time Transport Protocol (RTP) [6] traffic, tends to be more sensitive to delay and less sensitive to errors. In order to support these differing QoS constraints over a single physical link between two Internet nodes (e.g., routers or personal computers), many nodes insert traffic with different QoS requirements into different queues. Then, by servicing the queues based on the different QoS requirements, the node attempts to provide the QoS desired by the different types of traffic. For PPP sessions, multiple queues over a single physical link are supported using the PPP Multilink Protocol (MP) [7]. In this configuration each queue is carried by a different PPP link. This feature allows PPP to support differing QoS requirements. For instance, in the example shown in Fig. 4a the system has negotiated three PPP links.

Since radio link bandwidth is a limited resource, we should consider protocol overhead when choosing a protocol that will be carried over the radio link. PPP has been designed to minimize its own protocol overhead. In addition, it supports the compression of network-layer protocol headers such as TCP and IP/UDP/RTP header compression, further reducing the overhead of carrying user traffic over radio links.

It is typical to operate HDR with a received signal-to-noise ratio that results in a physical-layer PER of approximately 1 percent. This error rate is significantly higher than the error rate seen on most wireline networks. Since most network protocols and most network applications were designed assuming wireline error rates, the wireless link error rate needs to be reduced. The most straightforward method of reducing the error rate is for access terminals to operate at a higher signal-to-noise ratio regime. However, the increase in the signal-to-noise ratio required to reach wireline error rates results is a substantial decrease in overall throughput. A more efficient method for decreasing the error rate of the wireless link is obtained by implementing a form of automatic repeat request (ARQ). HDR implements a negative acknowledgment (NACK)-based radio link protocol (RLP) whereby incorrectly received blocks of data are detected and then retransmitted. This allows PPP and the higher layers to operate at an error rate regime similar to that experienced in wireline networks.

As shown in Fig. 4a, each PPP link may be carried by a separate RLP stream. In this specific example the system has negotiated three separate RLP streams to carry the three PPP links. This introduces the flexibility of allowing for finer control of the QoS. For instance, depending on the QoS requirement, different transmit scheduling policies with differing priorities may be implemented on some PPP streams. Additionally,

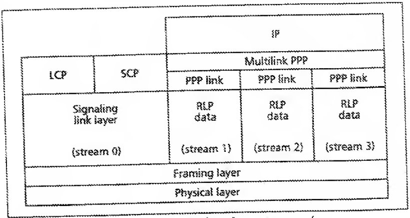

Figure 4a. *The air interface protocol stack — an example.*

RLPs with different effective error rates may be used on other PPP links. The framing layer shown in Fig. 4a is responsible for multiplexing the separate RLP streams into one physical layer.

In addition to user traffic, the HDR radio link must support the transport of signaling messages. The model for the transport of signaling streams is based on PPP. Signaling is partitioned into two basic types: the Link Control Protocol (LCP) and Stream Control Protocol (SCP). Similar to the PPP LCP, the LCP is used to negotiate radio link protocols and options at the start of the session and to control the radio link during the session. For example, the LCP is used at the start of the session to negotiate the link layer authentication type that will be used for the duration of the session. Similar to the PPP Network Control Protocol (NCP), the SCP is used to carry stream-specific signaling messages. For example, SCP is used to transmit the RLP NACKs upon detection of missing RLP data.

In the remainder of this section we discuss in which elements of the network the various layers of the protocols may be implemented. First we will describe the implementation on the user side of the air interface, to be followed by a brief description of the network side.

On the user side of the air interface reside two basic functional elements: the access terminal and the computer. These elements may reside in two devices, as in the case of a wireless HDR modem connected to a portable computer, or may be combined into a single device such as a wireless personal digital assistant (PDA). In the latter case, the device must implement the entire protocol stack, while in the former the protocol stack implementation may be partitioned in two ways.

In the first partitioning method, the access terminal implements the entire protocol stack. This partitioning is sometimes referred to as the *network model*. When using this partitioning, the access terminal and computer may physically be connected over Ethernet, through a PCMCIA interface, or over the Universal Serial Bus (USB). Figure 4b shows the layering endpoints of the network model.

In the second partitioning method, the access terminal implements the entire protocol stack with the exception of PPP and everything above PPP.

A first fundamental design choice is to separate the services, that is, low-rate data (voice being the primary service in this category) from high-rate data services, by using possibly adjacent but non-overlapping spectrum allocations.

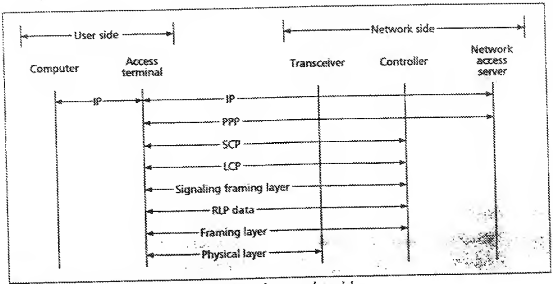

■ Figure 4b. *Air interface protocol endpoints — the network model.*

while the computer implements PPP and all protocols above PPP. This partitioning is sometimes referred to as the *relay model*. In the relay model, the access terminal and computer may be physically connected via RS-232 or USB. Figure 4c shows the layering endpoints of the relay model.

The network side of the air interface is modeled on the traditional Internet network access server (NAS) [8]. There are three basic functional elements: the transceiver, the controller, and the NAS. The transceiver implements the physical layer. The controller implements the framing layer, the RLP data layer, the signaling link layer, and all layers above the signaling link layer. The three functional blocks communicate over IP using open interfaces. The NAS implements the PPP layer and all layers above PPP. Figures 4b and 4c show the layering endpoints for the partitioning.

In our design of the access network, the access point implements the transceiver, the controller, and the NAS functional elements. The network interface implements the protocols and interfaces needed to connect the access point to an IP network and a backhaul network. Since the transceiver, controller, and NAS communicate over IP using open interfaces, there is no strict requirement for all the elements to be located in the access point. For example, in a more traditional cellular implementation, one might choose to centralize the controller and the NAS.

Only the transceiver and controller are specific to the radio link. The NAS and network interface are standard equipment used by today's Internet service providers (ISPs). By using an interface such as the widely supported Layer Two Tunneling Protocol (L2TP) between the NAS and the modem pool controller, it is possible to use this standard ISP equipment for many applications.

With the exception of other access points, the access point communicates with all elements in the access network using widely deployed Internet protocols. In addition, the access point transports all traffic using IP. Therefore, with the exception of the access point itself, all equipment in an access network is readily available Internet equipment.

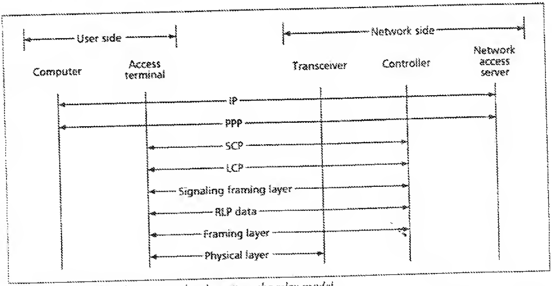

■ Figure 4c. *Air interface protocol endpoints — the relay model.*

ECONOMICS AND TARIFF CONSIDERATIONS

We consider finally the critical issue of the value of the service and how to establish tariffs. For truly nomadic users, constant travelers, people who prefer to work on their patio, at the beach or on the slopes, and so on, the service is most valuable and cannot be compared with high-speed wireline services provided by DSL or cable modems. At the other extreme for strictly fixed users, whose offices or homes are connected by fiber or wireline/cable services, the economics usually favor the latter. Suppose, however, that a carrier must decide between a wireline/cable high-speed solution or the wireless approach of this article. Here capital expenditures and possible tariff considerations dominate. The problem for wireless is that established wireline services have already conditioned users to expect a flat monthly rate essentially independent of the amount of service. The idealized economic model for digital packet-based wireless, practical considerations aside, will be to charge for usage on a packet basis.[7]

Most likely, however, most users will constitute a population of varying degrees of nomadicity, but even the occasional nomad will grow to depend on the flexibility and continuity provided by the "anywhere, anytime" nature of wireless connectivity. If this indeed turns out to be the case, even private networks (LANs) may be supplanted by wireless HDR usage. Economies of scale would seem to favor the HDR microcell located in or near a corporate campus over the private network. In fact, the best architecture for the latter may coincide with that of the former. Were this to occur, fiber, wire, and cable pipes may be relegated to the backbone and to the occasional supercomputer node, with the last kilometer or less becoming universally wireless for the majority of users.

Most users will likely constitute a population of varying degrees of nomadicity, but even the occasional nomads will grow to depend on the flexibility and continuity provided by the "anywhere, anytime" nature of wireless connectivity.

REFERENCES

[1] EIA/TIA/95, "Mobile Station-Base Station Compatibility Standard for Dual Mode Wideband Spread Spectrum Cellular System."

[2] S. Benedetto *et al.*, "Serial Concatenation of Interleaved Codes: Performance Analysis Design and Iterative Decoding," JPL TDA Progress Report 42-126, Aug. 15, 1996; see also *Proc. 1997 IEEE Int'l. Symp. Info. Theory*, Ulm, Germany, June 29–July 4, 1997, p. 106.

[3] "Internet Protocol," RFC 791, Sept. 1981.

[4] W. Simpson, "The Point-to-Point Protocol (PPP)," RFC 1661, July 1994.

[5] "Transmission Control Protocol," RFC 793, Sept. 1981.

[6] H. Schulzrinne *et al.*, "RTP: A Transport Protocol for Real-Time Applications," RFC 1889, Jan. 1996.

[7] K. Sklower *et al.*, "The PPP Multilink Protocol (MP)," RFC 1990, Aug. 1996.

[8] M. Beadles, "The Network Access Server," draft-beadles-nas-01.txt, Nov. 1998, work in progress.

BIOGRAPHY

PAUL BENDER is vice president of technology for Corporate Research and Development of QUALCOMM Incorporated. He received his Ph.D. in electrical engineering in 1992 with an emphasis on communications theory and systems at the University of California, San Diego (UCSD). He joined QUALCOMM in August 1992 and has contributed his expertise to the development of innovative CDMA digital communications at QUALCOMM. His responsibilities have included building the system test team and leading the system engineering team for the (IS-95) cellular and PCS infrastructure. He is currently involved in the development of QUALCOMM's High Data Rate (HDR) technology. Prior to joining QUALCOMM, he conducted research in signal processing for digital recording at UCSD's Center for Magnetic Recording Research. In addition, he assisted students in designing and debugging basic analog circuits for the Electrical and Computing Engineering Department at UCSD.

PETER J. BLACK is vice president of technology for Corporate Research and Development, of QUALCOMM Incorporated. He joined QUALCOMM in April 1993, where he was first engaged in the system design and development of CDMA mobile station ASICs. Since 1997 he has been working on next-generation system designs for high-speed mobile wireless packet data services as one of the HDR co-project engineers. Prior to joining QUALCOMM he worked at Austek Microsystems in Adelaide, Australia on full custom VLSI designs for digital signal processing. He received his B.E. degree in electrical engineering from the University of Queensland, Australia, in 1985. He received his M.S.E.E. and Ph.D. degrees from Stanford University, California, in 1990 and 1993, respectively. He is a Fulbright Scholar and was awarded the University Medal by the University of Queensland in 1985. Additionally, he has been awarded several patents throughout his career.

MATTHEW S. GROB is vice president, engineering for Corporate Research and Development, QUALCOMM Incorporated. He joined QUALCOMM in August 1991. Since 1997 he has been HDR co-project engineer and is head of Corporate Research and Development, System Engineering. He was the Globalstar Air Interface editor and co-author of the Globalstar data services specification. In 1995 he was the project engineer for the CONDOR proof-of-concept program which demonstrated secure CDMA and AMPS voice with point-to-point and net broadcast modes. His additional contributions to QUALCOMM include direct involvement in CDMA data standards IS-99, IS-657, and IS-707, as well as being awarded several patents. He holds a B.S.E.E. from Bradley University, Peoria, Illinois as well as an M.S.E.E. from Stanford University.

ROBERTO PADOVANI received a Laureate degree from the University of Padova, Italy and a M.S. and Ph.D. degrees from the University of Massachusetts, Amherst in 1978, 1983, and 1985, respectively, all in Electrical and Computer Engineering. In 1984 he joined M/A-COM Linkabit in San Diego where he was involved in the design and development of satellite communication systems, secure video systems, and error-correcting coding equipment. In 1986 he joined QUALCOMM, Inc. where he is now a Senior Vice President of Corporate Research and Development. Dr. Padovani has been involved in the design, development, and standardization of IS-95 based systems. He holds 35 patents on wireless CDMA systems. He is an IEEE Fellow, and was a co-recipient of the 1991 IEEE Vehicular Technology Society Best Paper Award for a fundamental paper on the capacity of a CDMA cellular system.

NAGABHUSHANA T. SINDHUSHAYANA earned his undergraduate degree in electrical engineering from the Indian Institute of Technology (IIT), Madras, India, in 1989. He pursued his graduate study at Cornell University, Ithaca, New York, where he obtained his Master's and Ph.D. degrees in 1992 and 1995, respectively. He did his graduate work in the area of coded modulation and symbolic dynamics, under the guidance of Prof. Chris Heegard. He has been with Qualcomm since 1994, where he has been working on turbo codes and wireless communications, including the HDR packet data system.

ANDREW J. VITERBI is a co-founder of QUALCOMM Incorporated. He has spent equal portions of his career in industry, having also co-founded a previous company, and in academia as a professor of engineering first at UCLA and then at UCSD, at which he is now Professor Emeritus. His principal research contribution, the Viterbi Algorithm, is used in most digital cellular phones and digital satellite receivers, as well as in such diverse fields as magnetic recording, speech recognition, and DNA sequence analysis. In recent years he has concentrated his efforts on establishing CDMA as the multiple access technology of choice for cellular telephony and wireless data communication. He has received numerous honors both in the United States and internationally. Among these are three honorary doctorates and memberships in both the National Academy of Engineering and the National Academy of Sciences. He currently serves on the President's Information Technology Advisory Committee.

[7] *The temptation to charge per slot allocated should be in any case resisted, for the disadvantaged user who requires more slots is already paying the price of increased latency. Furthermore, if low-rate users predominate in a certain area, the provider will have the incentive to improve service by allocating a new microcell.*